

ABSTRACT OF THE DISCLOSURE

A peer-to-peer storage system includes a storage coordinator that centrally manages distributed storage resources in accordance with system policies administered through a central administrative console. The storage resources, or "nodes," are otherwise unused portions of storage media, e.g., hard disks, that are included in the devices such as personal computers, workstations, laptops, file servers, and so forth, that are connected to a corporate computer network, and are thus otherwise available only individually to the respective devices. The storage coordinator assigns the nodes to various "replication groups" and allocates the storage resources on each of the nodes in a given group to maintaining dynamically replicated versions of the group files. The storage nodes in a given group perform dynamic file replication and synchronization operations by communicating directly, that is, peer-to-peer, using a message-based protocol. The storage coordinator also manages distributed searches of file content on the network by selecting one node from each group to search through the associated group files. The selected nodes report the search results back to the storage coordinator, which organizes the results and provides them to the user. Thereafter, in response to a request for various files by the user, the storage coordinator instructs the nodes that are near neighbors of the user to provide the requested files. The storage coordinator thus ensures that the amount of the network bandwidth consumed by the search operation is minimized.